

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 1 209 621 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
29.05.2002 Bulletin 2002/22

(51) Int Cl.⁷: G06T 5/20

(21) Application number: 01203319.7

(22) Date of filing: 03.09.2001

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 13.09.2000 US 661268

(71) Applicant: EASTMAN KODAK COMPANY
Rochester, New York 14650 (US)

(72) Inventors:
• Gallagher, Andrew C.,
c/o Eastman Kodak Company
Rochester, New York 14650-2201 (US)

• Gindele, Edward B.,
c/o Eastman Kodak Company
Rochester, New York 14650-2201 (US)

(74) Representative:
Nunney, Ronald Frederick Adolphe et al
Kodak Limited,
Patents, W92-3A,
Headstone Drive
Harrow, Middlesex HA1 4TY (GB)

(54) A method for enhancing a digital image based upon pixel color

(57) A method is described for enhancing a color digital image by first identifying a local neighborhood of pixels including a pixel of interest and then using the color values of one or more of the pixels from the local neighborhood to calculate a color weighting factor for the pixel of interest, where the color weighting factor is a near continuous function of the location of the pixel of interest in a color space. A spatial filter is applied to the

values of pixels from the local neighborhood to calculate a first signal value and a second signal value for the pixel of interest. The color weighting factor is then used to modify either the first value or the second value, and an enhanced pixel value is then generated for the pixel of interest from the first value and second values. The process is then repeated for other pixels in the color digital image.




FIG. 11

EP 1 209 621 A2

Description

FIELD OF THE INVENTION

- 5 [0001] The invention relates generally to the field of digital image processing and, more particularly, to a method for enhancing the texture of a digital image.

BACKGROUND OF THE INVENTION

- 10 [0002] Traditional methods of increasing the apparent sharpness of a digital image, such as the technique of unsharp masking, often produce unwanted artifacts at large transition edges in the image. For example, unsharp masking is often described by the equation:

$$15 \quad \text{Sproc} = \text{Sorg} + B(\text{Sorg} - \text{Sus})$$

where Sproc represents the processed image signal in which the high frequency components have been amplified, Sorg represents the original image signal, Sus represents the unsharp image signal, typically a smoothed image signal obtained by filtering the original image, and B represents the high frequency emphasis coefficient.

- 20 [0003] The unsharp masking operation may be modeled as a linear system. Thus, the magnitude of any frequency in Sproc is directly dependent upon the magnitude of that frequency in the Sorg image signal. As a consequence of this superposition principle, large edges in the Sorg image signal will often display a ringing artifact in the Sproc signal when the desired level of high frequency enhancement has been performed in other areas of the Sproc signal. This ringing artifact appears as a light or dark outline around the large edge, and may be visually objectionable.

- 25 [0004] Many non-linear filters based on local statistics exist for the purposes of noise reduction, sharpening, and contrast adjustment. For example, the median filter is well known in the art. In this filter, typically implemented for noise reduction, each pixel is replaced with the median value of some surrounding neighborhood. This filtering process is generally very successful at removing impulse noise; however, the processed image appears slightly less sharp than the original.

- 30 [0005] Another example of a non-linear filter based on local statistics is local histogram equalization, referred to as adaptive histogram modification by William Pratt on pages 278 - 284 of the book Digital Image Processing, Second Edition, John Wiley & Sons, 1991. With this filter, the values of pixels are modified by the cumulative histogram of a local window. This technique effectively adjusts the contrast of each region of a digital image, effectively increasing the local contrast in some regions of the image, and decreasing the contrast in other regions. This technique does not intend to increase the apparent sharpness of any given region. Also, this technique does not ensure that the typical artifacts of ringing will not occur.

- 35 [0006] There exist many algorithms for sharpening the appearance of images without generating artifacts or enhancing the notability of noise. In U.S. Patent 4,571,635, Mahmoodi and Nelson teach the method of deriving an emphasis coefficient B that is used to scale the high frequency information of the digital image depending upon the standard deviation of image pixel values in a local neighborhood. In addition, in U.S. Patent 5,081,692, Kwon and Liang teach that the emphasis coefficient B is based upon a center weighted variance calculation. However, neither Mahmoodi et al nor Kwon et al consider the expected standard deviation of noise inherent in the imaging system. By not considering the noise inherent in the imaging system, both Mahmoodi and Kwon make the implicit assumption that all imaging sources and intensities have the same noise characteristics. In addition, neither use separate strategies for texture and edge regions.

- 45 [0007] In U.S. Patent 4,794,531, Morishita et al teaches a method of generating the unsharp image with a filter whose weights on neighboring pixels are based upon the absolute difference between the central pixel and the neighboring pixel. Morishita claims that this method effectively reduces artifacts seen at the edges of a sharpened image (as compared with traditional unsharp masking). In addition, Morishita sets a gain parameter based upon local standard deviation and the standard deviation of the entire image. Again, Morishita does not consider the levels of noise inherent in the imaging system in order to approximate signal to noise ratios. In addition, the method of Morishita does not offer explicit control over the amount of edge reshaping.

- 50 [0008] In U.S. Patent 5,038,388, Song teaches a method of amplifying image details without amplifying the image noise by adaptively amplifying the high-frequency component of the image. An estimate of image noise power is used, however; this noise power is not described as being dependent on the intensity or the pixel. In addition, Song does not attempt to estimate signal to noise ratio in order to control the level of sharpening.

- 55 [0009] In US. Patent 4,689,666 Hatanaka discloses a method of using the color characteristics of a color digital image for the purposes of reducing the noise component of the color digital image. Hatanaka describes a process of

extracting color data for each picture element of the image, discriminating regions of the color digital image exhibiting a specific color on the basis of the extracted color data, and subjecting the image signal to spatial image processing for elimination of noise under different processing conditions for regions exhibiting the specific color and the remaining regions not exhibiting the specific color. Thus the method taught by Hatanaka has as a fundamental step the segmentation, or discrimination, of each pixel as belonging to the specific color or not belonging to the specific color. The step color discrimination can lead to unevenness the processed images due to the on/off nature of the color identification process.

[0010] In US. Patent 5,682,443 Gouch and MacDonald disclose a method of processing color digital images for the purpose of spatial sharpness characteristic enhancement. A method of unsharp masking is described which separates each color channel of the original color digital image into two parts based solely on the spatial frequency content of the original color channel. The difference between the original color channel and a low spatial frequency component of the original color channel forms a fringe component, or high spatial frequency component of the original color channel. Gouch and MacDonald teach a method of modifying the fringe component based on the color of either the low spatial frequency component or the original pixel values. The color considered is derived from the identical color channels that are sharpened. Their patent also discloses that the preferred method of implementing this feature uses a continuous mathematical function of color. The method disclosed by Gouch and MacDonald takes advantage of color as an image characteristic for enhancing the spatial detail. However, the unsharp masking procedure employed by Gouch and MacDonald has several shortcomings. First, their method only considers cases where each color channel of a color digital image undergoes the described unsharp masking operation. The variable weight used to modify the fringe data is derived from all of the color channels. Their method fails to consider cases when only a single digital image channel from a color digital image is sharpened, as is often the case when an image is represented as a luminance channel and a collection of chrominance channels. Additionally, Gouch and MacDonald necessitate that the modified fringe data is combined with the original pixel data. However, in many applications, it is preferred to modify the signal to which the modified fringe data is added. Thus, their method fails to consider cases where it is desirable to modify the signal to which the modified fringe data is combined.

[0011] Thus, there exists a need for an alternative method of manipulating a digital image in order to generate an image signal that appears to be sharper, or more in focus, while minimizing the ringing artifact that is evident with the unsharp masking technique and enhancing the magnitude of detail in the scene in a noise sensitive manner.

SUMMARY OF THE INVENTION

[0012] It is an object of this invention to allow for independent control of the enhancement applied to detail, large edges, and noisy regions.

[0013] The present invention is directed to overcoming one or more of the problems set forth above. Briefly summarized, according to one aspect of the present invention, the invention resides in a method for enhancing a color digital image by first identifying a local neighborhood of pixels including a pixel of interest and then using the color values of one or more of the pixels from the local neighborhood to calculate a color weighting factor for the pixel of interest, where the color weighting factor is a near continuous function of the location of the one or more local pixels from the local neighborhood in a color space. A spatial filter is applied to the values of pixels from the local neighborhood to calculate a first signal value and a second signal value for the pixel of interest. The color weighting factor is then used to modify either the first value or the second value, and an enhanced pixel value is then generated for the pixel of interest from a combination of the first value and second values. The process is then repeated for other pixels in the color digital image.

[0014] The present invention has the advantage of boosting the texture signal by a factor that is related to an estimate of the local color in the image. Thus, the portion of the texture signal coincident with regions of the digital image channel having certain local colors, such as the green of a grassy field due to many blades of grass in addition to image system noise, will experience a greater level of boost as compared with the portion of the texture signal associated with certain other local colors, such as a large region of clear blue sky where the only modulation is likely to be noise resulting from the imaging system. Therefore, while it is not desirable to increase the amplitude of regions having only noise modulation, it is preferable to boost the modulation where it can be attributed to actual modulation in the scene. In the present invention, the local color functions as a classifier that may be used to distinguish between, e.g., the two aforementioned types of regions in a scene.

[0015] These and other aspects, objects, features and advantages of the present invention will be more clearly understood and appreciated from a review of the following detailed description of the preferred embodiments and appended claims, and by reference to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016]

Fig. 1 is a block diagram illustrating an overview of the present method, including the invention as later described in more detail in connection with Figs. 11 and 15B.

Fig. 2 is a block diagram of the preference sharpener in Fig. 1 in more detail.

Fig. 3 is a block diagram of a first embodiment of the sharpening processor shown in Fig. 2.

Fig. 4 is a block diagram of a second embodiment of the sharpening processor shown in Fig. 2.

Fig. 5 is a block diagram of a first embodiment of the pedestal splitter shown in Fig. 3.

Fig. 5A is a block diagram of a second embodiment of the pedestal splitter shown in Fig. 3.

Fig. 5B is a block diagram of a third embodiment of the pedestal splitter shown in Fig. 3.

Fig. 6 is a block diagram of the avoidance signal generator shown in Fig. 5.

Fig. 7 is an example of an artifact avoidance function applied by the avoidance signal generator shown in Fig. 6.

Fig. 8 is a block diagram of the pedestal modifier shown in Fig. 3.

Fig. 9 is an example of a portion of an image showing a central pixel and the associated local region of the image.

Fig. 10 is an example of several tone scale functions.

Fig. 11 is a block diagram according to the invention of the texture modifier shown in Fig. 3.

Fig. 12 is a block diagram of a first embodiment of the local SNR estimator shown in Fig. 11.

Fig. 13 is a block diagram of a second embodiment of the local SNR estimator shown in Fig. 11.

Fig. 14 is a plot of a function used to generate a noise table.

Fig. 15A is a plot of a look up table (LUT) used by the SNR gain determiner shown in Fig. 11.

Fig. 15B is a plot of a look up table (LUT) used by the color gain determiner shown in Fig. 11.

DETAILED DESCRIPTION OF THE INVENTION

[0017] In the following description, the present invention will be described as a method implemented as a software program. Those skilled in the art will readily recognize that the equivalent of such software may also be constructed in hardware. Because image enhancement algorithms and methods are well known, the present description will be directed in particular to algorithm and method steps forming part of, or cooperating more directly with, the method in accordance with the present invention. Other parts of such algorithms and methods, and hardware and/or software for producing and otherwise processing the image signals, not specifically shown or described herein may be selected from such materials, components and elements known in the art. Given the system as described according to the invention in the following materials, software not specifically shown, suggested or described herein that is useful for implementation of the invention is conventional and within the ordinary skill in such arts.

[0018] It is instructive to note that the present invention utilizes a digital color image which is typically a two-dimensional array of red, green, and blue pixel values corresponding to light intensities. As used herein the term digital image refers to the color image, and in particular to a whole two dimensional array, or any portion thereof that is to be processed. In addition, the preferred embodiment is described with reference to an image of 1024 rows of pixels and 1536 lines of pixels, although those skilled in the art will recognize that digital images of different resolutions and dimensions may be used with equal, or at least acceptable, success. With regard to matters of nomenclature, the value of a pixel of a digital image located at coordinates (x,y) , referring to the x^{th} row and the y^{th} column of the digital image, shall herein comprise a triad of values $[r(x,y), g(x,y), b(x,y)]$ respectively referring to the values of the red, green and blue digital image channels at location (x,y) . In this regard, a digital image may be considered as comprising a certain number of digital image channels. In the case of a digital image comprising red, green and blue two-dimensional arrays, the image comprises three channels, namely, red, green and blue channels. Additionally, a luminance channel n may be formed from the color channels. The value of a pixel of a digital image channel n located at coordinates (x,y) , referring to the x^{th} row and the y^{th} column of the digital image channel, shall herein be a single value referred to as $n(x,y)$.

[0019] Referring to the overview of the invention shown in Fig. 1, a digital image is input to a preference sharpener 2, which sharpens the input digital image for the preference of an operator. The purpose of the preference sharpener 2 is to enhance the detail present in the digital image without enhancing noise or creating artifacts. This purpose is accomplished by a process that decomposes a digital image channel into a signal corresponding to the image detail, a signal corresponding to the image edges, a signal corresponding to local color and a signal corresponding to the image signal to noise ratio (SNR), as will be described in detail hereinbelow. This decomposition and creation of separate signals allows for independent control of the enhancement applied to detail, large edges, and noisy regions. The output of the preference sharpener 2 comprises a digital image that, in accordance with the invention, appears sharper and more natural than can be achieved with traditional means of sharpening a digital image. The particular approach of this invention, which relates to the independent control of the enhancement applied to detail based on local color,

will be described in detail in relation to Figures 11 and 15B after other elements of the preferred embodiment have been described in relation to the intervening Figures.

[0020] Fig. 2 shows a block diagram of the preference sharpener 2. The digital color image and noise information (in the form of a noise table 6) are input to the preference sharpener 2. The noise information is input in the form of a luminance table. The digital image is input to a luminance/chrominance converter 10 in order to generate a luminance digital image channel $n(x,y)$ and two chrominance digital image channels $gm(x,y)$ and $ill(x,y)$. The luminance digital image channel $n(x,y)$ and the luminance noise table are input over respective lines 11a and 11b to a sharpening processor 20 for enhancing the digital image channel $n(x,y)$ based on knowledge of the system noise. The transformation performed in the converter 10 from the digital color image, typically in an RGB space comprising separate red, green and blue image channels, to a luminance-chrominance color space is done by means of a color space matrix transformation, which results in the luminance digital image channel $n(x,y)$ and the two chrominance digital image channels $gm(x,y)$ and $ill(x,y)$, as is well known in the art. In accordance with the invention, as will be explained later, the two chrominance digital image channels $gm(x,y)$ and $ill(x,y)$ are also input to the sharpening processor 20. In the preferred embodiment, the matrix transformation to the luminance-chrominance space may be described with the following equation:

$$\begin{bmatrix} n(i, j) \\ gm(i, j) \\ ill(i, j) \end{bmatrix} = \begin{bmatrix} 1/3 & 1/3 & 1/3 \\ -1/4 & 1/2 & -1/4 \\ -1/2 & 0 & 1/2 \end{bmatrix} \begin{bmatrix} r(i, j) \\ g(i, j) \\ b(i, j) \end{bmatrix} \quad (A1)$$

Equation A1 provides a matrix rotation from RGB space into a luminance-chrominance space in order to isolate a single channel upon which the tone scale function (to be described) can operate. For example, a film scanner records (for each pixel) the red, green and blue printing density at that point. These three values can be thought of as a pixel's location in three-dimensional space. The converter 10 performs an axis rotation as shown in equation A1, which provides for a neutral axis, upon which $R = G = B$, and two color difference axes, green-magenta and illuminant. Note that rather than using the matrix rotation specified with Equation A1, standard color space descriptions such as CIE $L^*A^*B^*$ or CIE $L^*C^*H^*$ may alternatively be determined by the luminance/chrominance converter 10. These color spaces are well known in the art of image processing.

[0021] The noise table input to the sharpening processor 20 provides the relationship between the signal intensity level i and the expected amount of noise $\sigma_N(i)$ for that intensity level. In the preferred embodiment, as will be described in more detail, the noise table is a two column table, wherein the first column represents an intensity level i and the second column represents the expected standard deviation of noise $\sigma_n(i)$ for that intensity level.

[0022] The chrominance digital image channels $gm(x,y)$ and $ill(x,y)$ are input over a line 11c to a chrominance processor 40 and may be adjusted as desired. For instance, the action of the chrominance processor 40 may be to scale the chrominance image channels by a constant greater than 1.0 in order to increase their apparent saturation. The operation of the chrominance processor is not of particular relevance to the present invention, and consequently it will not be discussed further, except to note that, in the present embodiment, the output of the chrominance processor 40 is maintained identical to its input.

[0023] The digital image channels output from the sharpening processor 20 and the digital image channels output from the chrominance processor 40 are input to an RGB converter 30 for conversion back to a digital image composed of a red, green, and blue digital image channels. This conversion is again accomplished with a matrix rotation (i.e., the inverse of the previous color rotation matrix performed by the converter 10). Inverting a 3 by 3 matrix is well known in the art and will not be further discussed. The output of the RGB converter 30 is a digital image that has been sharpened for the preference of the operator.

[0024] As described in the preferred embodiment, the sharpening processor 20 operates upon only a luminance digital image channel. However, as an alternative embodiment, the sharpening processor could be applied to each of the red, green, and blue digital image channels. In this case, the digital color image signals (which are RGB signals) would be directly applied to the sharpening processor 20.

[0025] In Fig. 3, which illustrates a block diagram of the sharpening processor 20, the digital image channel $n(x,y)$ is shown to be separated into two portions by a pedestal splitter 50, that is, into a pedestal signal and a texture signal. The preferred embodiment of the pedestal splitter 50, shown in Fig. 5, outputs a texture signal $n_{tx}(x,y)$ and a pedestal signal $n_{ped}(x,y)$. The texture signal $n_{tx}(x,y)$ primarily contains the detail of the image, as well as the image noise. The pedestal signal $n_{ped}(x,y)$ is a conceptually smooth signal, except at regions of the image containing large occlusion boundaries, where the pedestal signal contains the transition edge. In the preferred embodiment, the sum of the texture and the pedestal signals is the digital image channel. For this reason, the digital image channel is said to have been

separated into two channels which together essentially comprise the digital image channel. An alternative embodiment of the pedestal splitter 50, shown in Figure 5A, essentially comprises a linear FIR filter. In this alternative embodiment, the pedestal signal $n_{ped}(x,y)$ is essentially equivalent to a lowpass signal and the texture signal $n_{txt}(x,y)$ is essentially equivalent to a highpass signal. Since other spatial frequency divisions are possible, those skilled in the art will understand that the present invention will yield beneficial results for any number of definitions of texture signals and pedestal signals. For example, Figure 5B describes an alternative pedestal splitter which utilizes a non-linear filter in order to separate a digital image channel into a pedestal signal $n_{ped}(x,y)$ and a texture signal $n_{txt}(x,y)$.

[0026] Referring again to Fig. 3, the digital image channel $n(x,y)$, the texture signal $n_{txt}(x,y)$ and the noise table 6 are input to a texture modifier 70. The purpose of the texture modifier 70 is to enhance the magnitude of detail in the scene in a noise sensitive manner. An estimate of the local signal/noise ratio (SNR) is made using information about the expected standard deviation of noise for a given intensity level as supplied by the noise table. This estimate of the local SNR is used to set the boost factor relating to the level of enhancement given to a local level of the texture signal. This procedure will be described in detail hereinbelow. The output of the texture modifier 70 is an enhanced texture signal $n'_{txt}(x,y)$. Furthermore, the pedestal signal $n_{ped}(x,y)$ is input to a pedestal modifier 60 for the purpose of increasing edge contrast to make the edges appear to have greater clarity and sharpness. The purpose of the pedestal modifier 60 is to enhance the image edges without producing artifacts. The method employed by the pedestal modifier 60 will be described in detail hereinbelow. The output of the pedestal modifier is an enhanced pedestal signal $n'_{ped}(x,y)$.

[0027] The outputs of the pedestal modifier 60 and the texture modifier 70 are input to an adder 80, which adds the enhanced texture signal $n'_{txt}(x,y)$ and the enhanced pedestal signal $n'_{ped}(x,y)$ together in order to produce the digital image channel output from the sharpening processor 20. The enhanced digital image channel $n'(x,y)$ may be expressed as

$$n'(x,y) = n'_{txt}(x,y) + n'_{ped}(x,y)$$

[0028] Fig. 4 shows an alternative embodiment of the sharpening processor 20, illustrating that minor perturbations to the arrangement of the components often have only a slight effect on the output of the sharpening processor 20. In this regard, the pedestal splitter 50 outputs the texture signal $n_{txt}(x,y)$ and an avoidance signal $a(x,y)$. The avoidance signal $a(x,y)$ is an intermediate signal calculated by the pedestal splitter 50 for determining the location of edges in the image. The derivation of this signal will be described below. The value of the avoidance signal ranges from 0.0 to 1.0. Image locations (x,y) having a value of the avoidance signal $a(x,y) = 0.0$ correspond with edge regions of the image; conversely, image locations (x,y) having a value of the avoidance signal $a(x,y) > 0.0$ correspond with regions of image detail or noise. As mentioned before, it is an object of this invention to allow for independent control of the enhancement applied to detail, large edges, and noisy regions. Thus, in the digital image channel output from the sharpening processor 20, it is desired that the edges (where $a(x,y) = 0$) be enhanced by the pedestal modifier 60, and that the regions of image detail or noise (corresponding to those locations where $a(x,y) > 0.0$) be enhanced by the texture modifier 70.

[0029] To this end, the digital image channel $n(x,y)$ is input to the pedestal modifier 60 and as before the texture signal $n_{txt}(x,y)$ is input to the texture modifier 70. The two resulting outputs from the pedestal modifier 60 and the texture modifier 70, the enhanced texture signal $n'_{txt}(x,y)$ and the enhanced pedestal signal $n'_{ped}(x,y)$, are input to an avoidance adder 81. The avoidance adder 81 requires three inputs: the two signals $n'_{txt}(x,y)$ and $n'_{ped}(x,y)$ to be summed, and the artifact avoidance signal $a(x,y)$. The two signals to be summed undergo a transformation such that one of the signals to be summed is multiplied by $a(x,y)$, and the other is multiplied by $(1-a(x,y))$. The two transformed signals are then added. The signal input which is scaled by $a(x,y)$ is known as the " $a(x,y)$ input" of the avoidance adder 81 and the signal input which is scaled by $(1-a(x,y))$ is known as the " $(1-a(x,y))$ " input of the avoidance adder 81. In the case of the present embodiment of the sharpening processor 20, the signal output from the avoidance adder can be expressed as:

$$n'(x,y) = a(x,y)n'_{txt}(x,y) + (1-a(x,y))n'_{ped}(x,y)$$

[0030] Referring to Fig. 5, the digital image channel $n(x,y)$ input into the pedestal splitter 50 is split into a high pass signal $n_{hp}(x,y)$ and a low pass signal $n_{lp}(x,y)$ by a frequency splitter 94. Although there are many known techniques for generating highpass and lowpass signals, the frequency splitter is preferably implemented with a Gaussian filter having a standard deviation (sigma) of 0.9 pixels. The preferred value of the standard deviation of the Gaussian filter may vary with image size. The value of 0.9 pixels for the standard deviation was derived by optimizing the present invention with 1024 by 1536 pixel size images. This Gaussian filter is a two dimensional circularly symmetric low-pass filter whose filter coefficients may be derived by the following formula which is well known in the art:

$$g(i,j) = 1/(\sigma \sqrt{2\pi}) \exp[-(i^2+j^2)/(2\sigma^2)]$$

Where

$g(i,j)$ = the Gaussian filter coefficient at the (i,j) th pixel
 σ = the standard deviation of the Gaussian filter (0.9)
 π = the constant approximately 3.1415...

In the preferred embodiment, i and j range from -3 to +3, inclusive, for a total of 49 coefficients for the filter $g(i,j)$. Note that another technique, well known in the art, exists for separating a Gaussian filter into horizontal and vertical components for subsequent application for the purpose of computational cost reduction. In either case, the frequency splitter 94 applies the Gaussian filter $g(i,j)$ to the digital image channel $n(x,y)$ by the process of convolution, which is well known in the art. The convolution may be expressed as:

$$n_{lp}(x,y) = \sum \sum n(x-i,y-j)g(i,j)$$

with the summations occurring over all i and j . The signal $n_{lp}(x,y)$ resulting from this convolution process is the low-pass signal. The low-pass signal is output from the frequency splitter 94. In addition, the high-pass signal $n_{hp}(x,y)$ is output from the frequency splitter 94 after having been derived by the following relationship:

$$n_{hp}(x,y) = n(x,y) - n_{lp}(x,y)$$

[0031] The low-pass signal $n_{lp}(x,y)$ is input into an avoidance signal generator 104 for forming the avoidance signal $a(x,y)$, as will be described in detail below. A texture signal generator 90 receives both the high-pass signal $n_{hp}(x,y)$ and the avoidance signal $a(x,y)$ and both signals are multiplied therein for producing the texture signal $n_{txt}(x,y)$. Thus, the texture signal $n_{txt}(x,y)$ may be expressed as:

$$n_{txt}(x,y) = a(x,y) * n_{hp}(x,y).$$

The texture signal $n_{txt}(x,y)$ computed by the texture signal generator 90 is then output by the pedestal splitter 50. In addition, as shown in broken line in Fig. 5, the avoidance signal $a(x,y)$ may optionally be output by the pedestal splitter 50, in particular to provide an input to the avoidance adder 81 in the alternative embodiment of the sharpening processor 20 shown in Fig. 4.

[0032] A pedestal generator 100 receives the luminance digital image channel $n(x,y)$ and the texture signal $n_{txt}(x,y)$ and subtracts the texture signal from the luminance signal, which produces the pedestal signal $n_{ped}(x,y)$. Thus, the pedestal signal may be expressed as:

$$n_{ped}(x,y) = n(x,y) - n_{txt}(x,y).$$

The pedestal signal $n_{ped}(x,y)$ computed by the pedestal signal generator 100 is then output from the pedestal splitter 50.

[0033] An alternative embodiment of the pedestal splitter 50 (identified as 50A) is shown in Fig. 5A. The luminance digital image channel $n(x,y)$ input into the alternative embodiment of the pedestal splitter 50A is split into a high pass signal and a low pass signal by a frequency splitter 94 preferably with the Gaussian filter as described above having a standard deviation (σ) of 0.9 pixels. The preferred value of the standard deviation of the Gaussian filter may vary with image size. The value of 0.9 pixels for the standard deviation was derived by optimizing the present invention with 1024 by 1536 pixel size images. The frequency splitter 94 applies the Gaussian filter $g(i,j)$ to the digital image channel $n(x,y)$ by the aforementioned process of convolution.

[0034] The signal $n_{lp}(x,y)$ resulting from this convolution process is the low-pass signal output from the frequency splitter 94. In addition, a high-pass signal $n_{hp}(x,y)$ is output from the frequency splitter 94 after having been derived by the following relationship:

$$n_{hp}(x,y) = n(x,y) - n_{lp}(x,y)$$

[0035] In this alternative embodiment of the pedestal splitter, the lowpass signal $n_{lp}(x,y)$ is output from the pedestal splitter as the pedestal signal $n_{ped}(x,y)$. (i.e., in this case, the pedestal signal $n_{ped}(x,y)$ is set equal to the lowpass signal $n_{lp}(x,y)$ rather than being derived with a control signal. This is equivalent to $a(x,y) = 1$.) Likewise, the texture signal $n_{bt}(x,y)$ is set equal to the highpass signal $n_{hp}(x,y)$.

[0036] This alternative embodiment of the pedestal splitter 50A simply performs a digital filtering operation by convolution and outputs the highpass and lowpass signals as the texture and pedestal signals, respectively. This embodiment of the pedestal splitter 50A is simpler than the preferred embodiment 50 because no control signal is used or calculated. However, the quality of the digital image output from the preference sharpener 4 employing this alternative embodiment is inferior to the quality of the digital image output from the preference sharpener 4 employing the preferred embodiment of the pedestal splitter 50.

[0037] Those skilled in the art will recognize that there are many types of operations that may be used to generate two signals that essentially sum to generate the digital image channel. The preferred embodiment of the pedestal splitter and the alternative embodiment of the pedestal splitter are both examples of operations that may be used to separate a digital image channel into two signals that essentially sum to generate the digital image channel.

[0038] Referring to Fig. 6, there is illustrated a block diagram of the avoidance signal generator 104. In this regard, the low pass signal $n_{lp}(x,y)$ is input into a non-directional squared gradient calculator 106 for producing a non-directional gradient signal. This calculation is performed by first calculating the difference between the pixel and its upper vertical neighbor, as well as the difference between the pixel and its horizontal neighbor to the right. The non-directional squared gradient is the sum of the squares of these two differences. The non-directional squared gradient signal, $ndg(x,y)$ may be expressed by the following equation:

$$ndg(x,y) = [n_{lp}(x,y) - n_{lp}(x-1,y)]^2 + [n_{lp}(x,y) - n_{lp}(x,y+1)]^2$$

The value of $ndg(x,y)$ is then mapped by an artifact avoidance function applicator 108 to produce an artifact avoidance signal $a(x,y)$ as the output of the avoidance signal generator 104.

[0039] Referring to Fig. 6, this mapping is performed by passing the non-directional squared gradient signal, $ndg(x,y)$, through an artifact avoidance function $av(y)$ formed in the following manner. The preferred embodiment utilizes an artifact avoidance function defined by:

$$av(y) = (1/2) (1 + \cos(\pi(y - C_0)/(C_1 - C_0)))$$

for $y > C_0$ and $y < C_1$,

$$av(y) = 0$$

for $y \geq C_1$ and

$$av(y) = 1$$

for $y \leq C_0$ where C_0 and C_1 are numerical constants.

[0040] The preferred values for C_0 and C_1 will vary according to the range of the input data. The input data range in the images used in the development of this preferred embodiment ranged from 0 to 4095. In this case, the preferred value for C_0 is 996, and the preferred value for C_1 is 8400. As an alternate method for constructing the artifact avoidance function is described in the following equation:

$$av(y) = 1 - \sqrt{(y - C_0) / (C_1 - C_0)}$$

for $y \geq C_0$ and $y \leq C_1$,

$$av(y) = 1$$

for $y < C_0$, and

$$av(y) = 0$$

for $y > C_1$. Fig. 7 shows an example plot of the function $av(y)$.

[0041] Note that in the preferred embodiment, the value of C_1 is related to the size sigma (in standard deviations) of the Gaussian filter. The value of C_1 is inversely proportional to the square of sigma (in standard deviations) of the Gaussian filter. Preferably, C_1 may be determined by the relation:

$$C_1 = 6804/(\text{sigma} * \text{sigma})$$

Also, C_0 is related to C_1 by the equation:

$$C_0 = 0.127 C_1^{-18}$$

[0042] The artifact avoidance signal, $a(x,y)$, is generated by the artifact avoidance function applicator 108. This is accomplished by applying the artifact avoidance function $av(y)$, to the non-directional gradient signal $ndg(x,y)$, as shown in Figure 7. The mathematical result is described by the equation:

$$a(x,y) = av(ndg(x,y)).$$

The artifact avoidance function shown in Figure 7 is most efficiently implemented as a look-up-table (LUT).

[0043] It is helpful to understand that the avoidance signal $a(x,y)$ is an example of a control signal created from a spatially filtered version of the digital image channel. To begin with, a generalized linear spatial filter is described by the equation

$$c(x,y) = \sum \sum d(x-i,y-j)g(i,j)$$

where the $d(x-i,y-j)$ values represent local pixel values surrounding the $(x,y)^{\text{th}}$ pixel and the $g(i,j)$ values represent numerical coefficients of a digital filter that do not depend on the pixel values, and are independent of x and y . A non-linear spatial filter herein described shall be defined as a spatial filter which cannot be described by the linear spatial filter equation. The output of applying a generalized control signal to an input signal bears a multiplicative relationship to the input signal. An example of applying a control signal is given by the equation:

$$c(x,y) = h(x,y)d(x,y)$$

where the $d(x,y)$ values represent the input signal $(x,y)^{\text{th}}$ pixel values and the $h(x,y)$ values represent the $(x,y)^{\text{th}}$ pixel values of the control signal. The result of applying a control signal to an input signal falls into the general category of a non-linear spatial filter if the control signal is derived from a spatially filtered version of the input signal. The avoidance signal $a(x,y)$ is an example of a control signal created from a spatially filtered version of the digital image channel. The texture signal $n(x,y)$ described in the preferred embodiment is an example of a non-linear spatial filter produced with the application of a control signal and a digital filter applied to the image channel $n(x,y)$.

[0044] Additionally, several other examples of non-linear filters which are well known in the art of digital image processing occur when the non-linear filter coefficients $g(i,j)$ are not independent of position x and y . More specifically, it is often the case that the non-linear filter coefficients $g(i,j)$ are dependent upon the value of pixels of $d(x,y)$ within a local neighborhood of (x,y) . In the preferred embodiment, the sum of the coefficients of $g(i,j)$ is 1.0, although this should not be seen as a limitation to the scope of the embodiment. Note that the median filter which is well known in the art of image processing is an example of a non-linear filter. In addition, the sigma filter described by Jong-Sen Lee in the article "Digital Image Enhancement and Noise Filtering by use of Local Statistics," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-2, No. 2, March 1980, is also an example of a non-linear filter. Note that in this description the non-linear filter is described as being only a function of the image signal $d(x,y)$, and thus is a Finite Impulse Response (FIR) filter. The non-linear filter may similarly be a recursive filter for which the coefficients of the recursive filter are dependent upon location (x,y) . For example, in commonly-assigned copending Patent Application Serial No. 09/457,036, filed December 8, 1999 and entitled "Adjusting the Contrast of a Digital Image with an Adaptive

Recursive Filter", therein is described an adaptive recursive filter for which the filter coefficients vary depending upon the local pixel values in a neighborhood of (x,y) . Fig 5B describes an alternative pedestal splitter which utilizes a non-linear filter in order to separate a digital image channel into a pedestal signal $n_{ped}(x,y)$ and a texture signal $n_{txt}(x,y)$. In this alternative pedestal splitter, the digital image channel $n(x,y)$ is input to the pedestal generator 51. The pedestal generator 51 calculates a pedestal signal $n_{ped}(x,y)$ from the digital image channel $n(x,y)$ by application of a non-linear filter to the digital image channel. The texture signal $n_{txt}(x,y)$ is then calculated by passing the negative of the pedestal signal $n_{ped}(x,y)$ and the digital image channel to an adder 53, as shown. The adder 53 determines the texture signal $n_{txt}(x,y)$ by applying the following formula:

$$n_{txt}(x,y) = n(x,y) - n_{ped}(x,y).$$

In this manner, the pedestal splitter 50B generates a pedestal signal $n_{ped}(x,y)$ and a texture signal $n_{txt}(x,y)$. The pedestal signal $n_{ped}(x,y)$ and the texture signal $n_{txt}(x,y)$ are output from the pedestal splitter 50B.

[0045] Referring now to Fig. 8, the pedestal modifier 60 shown therein is a sharpening filter algorithm that is adaptive to the morphology of a digital image channel, as based on one or more image characteristics calculated over a localized region of the digital image channel (as isolated by a two-dimensional window). The filtering algorithm includes a stage for normalizing the tone scale conversion with a scaling function for the image statistic(s) over the localized region of a digital image channel and then using the scaled tone scale to reshape the structure of edges in the digital image channel. More specifically, the shape of the local tone scale function is ordinarily selected before processing the image, but the exact scaling and translation of the local tone scale operation is determined from the statistics in any localized region. Accordingly, after the region is isolated, implementation of the algorithm includes identifying a desired statistical characteristic from the local region, performing a local tone scale conversion that is normalized in relation to the statistical characteristic, and mapping the central pixel of the local region through the tone scale conversion to generate an enhanced central pixel value. More detail of the pedestal modifier 60 is described in commonly assigned U.S. Patent Application Serial No. 09/324,239, entitled "A Method for Enhancing the Edge Contrast of a Digital Image" and filed June 2, 1999 in the names of A.G. Gallagher and E.B. Gindele, and which is incorporated herein by reference.

[0046] The pedestal modifier 60 may be implemented in several modes. More specifically, and particularly as shown in the block diagram of Figure 8, the pedestal modifier 60 is preferably practiced in a first mode by scaling the input pixel values to the local maximum and minimum for each region and then applying the scaled input values to a tone scale function with normalized coordinates; the output obtained by this conversion operation is then processed with an inverse of the scaling function (used to generate the input) to obtain the enhanced output value. Alternatively, in a second mode of implementing the pedestal modifier 60, a local tone scale is constructed having coordinates that are bounded by the local maximum and minimum for each region, and then each central pixel value is mapped through the bounded coordinates of the tone scale function. This will directly produce the enhanced output value. In a third mode, only the input of the tone scale function is scaled with the inverse scaling function and the output pixel is processed with the inverse of the scaling function; in a fourth mode, only the output of the tone scale function is scaled and the input pixel is processed with the scaling function. In any mode, the shape of the tone scale function may remain fixed over the image; however, it may be desirable to vary the shape of the tone scale function dependent upon the statistical characteristics of the region. Furthermore, while it may be typical to apply the normalized tone scale operation to substantially all the pixels in the input digital image channel, it may be desirable to only apply the method to selected areas of the digital image channel. For details of the second, third and fourth modes, the reader is directed to the aforementioned Serial No. 09/324,239, which is incorporated herein by reference. These alternative modes provide a processed digital image channel identical to the processed digital image channel obtained by the pedestal modifier 60 disclosed in the preferred embodiment (i.e., the first mode, Fig. 8).

[0047] Referring further to Fig. 8, there is illustrated a block diagram of a preferred embodiment for implementing the pedestal modifier 60 according to the first mode as set forth above. It also facilitates understanding to note that the preferred embodiment of the pedestal modifier 60 is described in reference to a relatively high resolution digital image channel having, e.g., dimensions of 1024 pixels in height by 1536 pixels in width, although the invention is similarly operative with image channels having greater or lesser image resolutions. The choice of image resolution, however, may have a concomitant effect on the window size, i.e., while a 5 X 5 window has been found preferable for 1024 X 1536 resolution, the choice of a higher resolution image sensor would ordinarily allow a larger window size, and vice versa for a lower resolution sensor.

[0048] With reference also to Fig. 9, the value of each input central pixel 110 of the digital representation of the digital image channel is first passed to a local statistics determiner 116 for calculating local statistics. The values of surrounding pixels 114 included within a local region 112 are also passed to the local statistics determiner 116. The local region 112 and the input pixel 110 constitute a spatial region of the digital image channel. In the preferred embodiment, the local region about each input pixel 110 is defined as a square neighborhood centered on the input pixel for which the

local statistics are being determined. In the preferred embodiment, the local region is a square neighborhood with 5 pixels on each side. Figure 9 illustrates the local region implemented in the preferred embodiment. The pixels 114 included within the local region 112 are illustrated as the 24 pixels (indicated by hatching) surrounding the input pixel 110 passed to the local statistics determiner 116. Those skilled in the art will recognize that the present invention will be operable with a large variety of local region sizes and shapes, and the usage (in this description and as claimed) of the term "central" to describe the processed pixel merely refers to a pixel having a local surround of pixels, rather than a pixel at the centroid of its local region.

[0049] In Figure 8, the value u of each central pixel 110 is input to the local statistics determiner 116 and a scaling function evaluator 120. In addition, the values of the pixels 114 in the local region are input to the local statistics determiner 116 in order to generate a statistical characteristic of the local region. A scaling function generator 118 uses the statistical characteristic to generate a scaling function $s(x)$ that is applied to the scaling function evaluator 120 in order to normalize the value u of each input pixel 110 and thereby generate a first intermediate value I . The statistical characteristic is also used by an inverse scaling function generator 122 to generate a function $s^{-1}(x)$ that is the inverse of the scaling function $s(x)$. The first intermediate value I is then applied to a tone scale applicator 124, which maps the intermediate value I through a tone scale function $f(x)$ provided by a tone scale function generator 126 in order to generate second intermediate value I_2 . The output intermediate value I_2 from the tone scale applicator 124 is then applied to an inverse scaling function applicator 128, which generates a processed value u' in the system image space by using the inverse function $s^{-1}(x)$ provided by the inverse scaling function generator 122.

[0050] More specifically, the local statistics determiner 116 calculates several characteristics that describe the image data contained in the local region 112 and the input pixel 110. The local statistics could be any mathematical combination of the pixel values contained in the local region 112 and the input pixel 110. In the preferred embodiment, the local statistics determiner 116 determines two local statistics, the maximum and minimum pixel values from the local region 112 and the input pixel 110. As an alternative, the local statistics determiner 116 could determine the mean of the n maximum pixel values, and the mean of the m minimum pixel values, where n and m are integers much smaller than the number of pixels contained in the local region 112, for example, $n = m = 3$. As will be described, these local statistics are utilized in the generation of the scaling function $s(x)$ and the inverse scaling function $s^{-1}(x)$.

[0051] The local statistics determined by the local statistics determiner 116 are then passed to the scaling function generator 118 and the inverse scaling function generator 122. In the preferred embodiment, the scaling function may be represented with the following equation:

$$s(u) = \frac{u - N - \text{min}}{X - N - \text{min}}$$

input value
min

where X represents the maximum pixel value from the local region 112 and N represents the minimum pixel value from the local region 112. The purpose of the scaling function is to modify the value of the central pixel 110 such that the modified pixel value may be used as an input to a tone scale function 130, as shown in Figure 10. In the preferred embodiment, this tone scale function 130 is applied as a look-up-table (LUT). In the preferred embodiment, the effect of the scaling function is to normalize (i.e., compress) the interval $[N X]$ of input values to an interval $[0 1]$, as shown in Fig. 10.

[0052] In a similar manner, the inverse scaling function generator 122 creates the inverse scaling function $s^{-1}(u)$; the inverse scaling function $s^{-1}(u)$ may be represented by the following equation:

$$s^{-1}(u) = u(X - N) + N$$

The purpose of the inverse scaling function is to scale the output values obtained from the tone scale function 130 back to the interval $[N X]$. The scaling function and the inverse scaling function are indeed inverses, as it can easily be shown that $s^{-1}(s(u)) = u$.

[0053] The value of the input pixel 110 is then passed to the scaling function evaluator 120 in order to generate an intermediate value $I = s(u)$, where I is the intermediate value, $s(x)$ is the scaling function, and u is the value of the input pixel 110. The output of the scaling function evaluator 120 is the intermediate value I . The intermediate value I is then passed to the tone scale applicator 124. The purpose of the tone scale applicator 124 is to affect the contrast of the value of the input pixel 110 with respect to the local region 112. The tone scale applicator 124 modifies the intermediate value I , thereby creating the second intermediate value I_2 . The tone scale applicator 124 modifies the intermediate value I with the tone scale function 130 (i.e., $f(x)$). Thus, the operation of the tone scale applicator may be expressed in the following equation: $I_2 = f(I)$, where I_2 is the second intermediate value, $f(x)$ is the tone scale function, and I is the intermediate value output from the scaling function evaluator 120.

[0054] The tone scale function $f(x)$ is generated by the tone scale function generator 126. In the preferred embodi-

ment, the tone scale function $f(x)$ is the sigmoid function, formed by integrating a Gaussian probability distribution function, and is well known in the art. The sigmoid function is fully characterized by the standard deviation σ of the Gaussian probability function that is integrated. The tone scale function $f(x)$ may be expressed by the following formula:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^x \exp\left(\frac{-x^2}{2\sigma^2}\right) dx$$

In the preferred embodiment, the tone scale function is found by computing the summation of discrete samples of a Gaussian distribution, as shown in the following formula:

$$f(x) = \frac{\sum_{u=-a}^{u=a(2x-1)} \exp\left(\frac{-x^2}{2\sigma^2}\right)}{\sum_{u=-a}^{u=a} \exp\left(\frac{-x^2}{2\sigma^2}\right)}$$

where $2a$ is the number of samples in the LUT. In the preferred embodiment, $a = 50$.

[0055] Notice that the slope of the tone scale function $f(x)$ evaluated at $x = 0.5$ may be determined by the relationship:

$$\gamma = \frac{1}{\sigma\sqrt{2\pi}},$$

where γ is the slope of the sigmoid function evaluated at $x = 0.5$, and the approximate slope of the summation described above, and σ is the standard deviation of the Gaussian probability function. When the value of $\gamma > 1.0$, the operation sharpens the digital image channel by reshaping the local structure of edges. When the value of $\gamma < 1.0$, the operation tends to produce an image with a softer appearance than the input image. In the preferred embodiment $\gamma = 2.1$ and $\sigma = 0.19$. In the preferred embodiment, the tone scale function $f(x)$ is independent of the local statistics output from the local statistics determiner 116, and thus need be computed only one time per image channel, despite the number of pixels being processed.

[0056] Figure 10 illustrates a variety of tone scale functions, one of which is shown as the tone scale function 130, and the associated standard deviations of the Gaussian probability function σ . Those skilled in the art will recognize that a wide variety of functions in addition to the sigmoid function will accomplish the goal of the tone scale applicator 124.

[0057] The second intermediate value l_2 output from the tone scale applicator 124 is passed to the inverse scaling function applicator 128. In addition, the inverse scaling function $s^{-1}(x)$ output from the inverse scaling function generator 122 is passed to the inverse scaling function applicator 128. The purpose of the inverse scaling function applicator 128 is to generate the processed pixel value u' from the second intermediate value l_2 . The inverse scaling function applicator 128 generates the enhanced pixel value u' by evaluating the inverse scaling function:

$$u' = s^{-1}(l_2)$$

Substituting for l_2 and l ,

$$u' = s^{-1}(\{f(s(u))\})$$

where $s^{-1}(x)$ represents the inverse scaling function, $f(x)$ represents the tone scale function, $s(x)$ represents the scaling function, u represents the value of the input pixel 110, and u' represents the value of the enhanced pixel. The digital image channel output from the pedestal modifier 60 is an enhanced pedestal signal $n'_{\text{ped}}(x,y)$.

[0058] Typically, the process for generating a processed pixel value u' from the pedestal modifier 60 is repeated for every pixel of the digital image channel in order to generate a processed pixel value corresponding to each pixel in the digital image channel. These processed pixel values, when considered as a collection, form a processed digital image channel. Alternatively, the process for generating a processed pixel value u' may be repeated for a subset of every pixel of the digital image channel in order to generate a processed pixel value corresponding to a subset of pixels in the digital image channel. In the latter case, the subset, when considered as a collection, forms a processed image for a portion of the channel. In general, the parameters of the present algorithm may be adjusted in order to produce an enhanced digital image channel, with the appearance of greater edge contrast and more clarity than the digital image channel.

[0059] The pedestal modifier 60 requires definition of two parameters: the size of the local region (which may be described by a window size for the sharpening filter) and the shape of the tone scale conversion function (which may be described by the slope of the tone scale function). The algorithm tends to perform best when the window size is kept fairly small, e.g., no greater than say about 7X7 pixels for a 1024 X 1536 imager. Indeed, as the window size and slope of the local tone scale function increase, the image begins to have a "paint-by-number" artifact. Also, a large window tends to compress texture on either side of a large occlusion boundary type edge in an image because the maximum and the minimum tend to occur on opposite sides of the edge. This effect is reduced by keeping the window size relatively small. A window size of 5 X 5 or fewer pixels has been found generally preferable for producing satisfactory results in typical situations for a 1024 X 1536 resolution. As mentioned earlier, other suitable window sizes may be empirically determined without difficulty for other image resolutions; alternatively, a larger window may be acceptable in certain situations for a given resolution if the user is willing to accept increased artifacts in order to obtain the sharpening effect of the larger window.

[0060] The pedestal modifier as herein described has the advantage of controlling the value of the central pixel in accordance with a statistical characteristic of the region, e.g., driving the value of the central pixel either toward the local maximum or the local minimum of the region, except in the case where the central pixel is substantially midway between the local maximum and the local minimum. Consequently, edge transitions occur over a narrower range of pixels than in the input image, thus generating an image that appears to be sharper, or more in focus, than the original image. Moreover, since the output of the tone scale conversion is modified by the statistical characteristic, e.g., bounded by the local maximum and the local minimum of the region, systematic overshoot and undershoot at an edge boundary is diminished and the ringing artifact is not as noticeable.

[0061] A detailed view of the texture modifier 70 can be seen in Fig. 11 in accordance with the invention. The texture modifier 70 inputs the digital image channel $n(x,y)$, the texture signal $n_{\text{tex}}(x,y)$, the luminance noise table 6 and the chrominance channels, $gm(x,y)$ and $ill(x,y)$. The output of the texture modifier 70 is an enhanced texture signal $n'_{\text{tex}}(x,y)$. The purpose of the texture modifier 70 is to boost the texture signal by a factor that is related to an estimate of the local signal to noise (SNR) ratio and an estimate of the local color, as determined by $gm(x,y)$ and $ill(x,y)$.

[0062] Thus, disregarding local color, the portion of the texture signal coincident with regions of the digital image channel having high SNR will experience a greater level of boost as compared with the portion of the texture signal associated with low SNR regions of the digital image channel. For example, suppose that an image to be sharpened contains a large region of clear blue sky. The only modulation in this region is likely to be noise resulting from the imaging system. Therefore, it is not desirable to increase to amplitude of this modulation. On the other hand, modulation of a grassy field is a result of the modulation due to many blades of grass in addition to image system noise. In this case, it is preferable to boost the modulation since it can be attributed to actual modulation in the scene. In the present invention, the signal to noise ratio is estimated to act as a classifier that may be used to distinguish between, e.g., the two aforementioned types of regions in a scene.

[0063] Considering only color, it may be desirable for the texture signal $n'_{\text{tex}}(x,y)$ to be enhanced either more or less, depending on the local color about location (x,y) . For example, any detail occurring on a clear blue sky may simply be from film grain. Thus, it would be desirable to deemphasize the sharpening that is applied to pixels in the original color digital image corresponding to the color blue. This may be accomplished by lowering the gain which is applied to the texture signal $n_{\text{tex}}(x,y)$ for locations in the original color digital image with blue colors.

[0064] In this regard, the digital image channel $n(x,y)$, the texture signal $n_{\text{tex}}(x,y)$ and the luminance noise table 6 are input to a local SNR estimator 160. The purpose of the local SNR estimator 160 is to estimate the local signal to noise ratio which will later be used to determine the level of boost given to a specific pixel of the texture signal. In the preferred embodiment, the output of the local SNR estimator 160 is one estimate of the SNR for each pixel in the digital image

channel, represented by the symbol $SNR(x,y)$. However, it is also possible that the SNR may be calculated for only a subset of the pixels' locations of the digital image channel. The local SNR estimator 160 will be described in greater detail hereinbelow.

[0065] The output of the local SNR estimator 160, $SNR(x,y)$, is passed into an SNR gain determiner 134. The SNR gain determiner 134 is meant to determine an individual level of boost $B(x,y)$ to apply at each particular (x,y) location of the texture signal, given that an estimate of the local SNR is known. This is accomplished by the use of a look up table (LUT), for example the LUT shown as a plot in Fig. 15. In the preferred embodiment, this LUT is monotonically increasing. The LUT output (the gain) increases with the increase of the LUT input (the estimate of the SNR.)

[0066] The equation for the LUT shown in Fig 15 is given as:

$$\begin{aligned} B(x,y) &= sf_{\min} && \text{for } SNR(x,y) < SNR_{\min} \\ B(x,y) &= sf_{\min} + (sf_{\max} - sf_{\min})(SNR(x,y) - SNR_{\min}) / (SNR_{\max} - SNR_{\min}) && \text{for } SNR_{\min} < SNR(x,y) < SNR_{\max} \\ B(x,y) &= sf_{\max} && \text{for } SNR(x,y) > SNR_{\max} \end{aligned}$$

Where sf_{\max} is $\geq sf_{\min}$, and $SNR_{\max} \geq SNR_{\min}$. In the preferred embodiment, sf_{\max} is 3.5, sf_{\min} is 2.0, SNR_{\max} is 2.0 and SNR_{\min} is 0.4.

[0067] Additionally, the chrominance channels $gm(x,y)$ and $ill(x,y)$ are input to the color gain determiner 135. The color gain determiner 135 calculates an individual level of boost $B_c(x,y)$ to apply at each location (x,y) of the texture signal, dependent upon the local color as described by the chrominance channels $gm(x,y)$ and $ill(x,y)$ and the distance of location (x,y) from identified color region(s) in a chrominance color space. The local color can be calculated from the pixel of interest, e.g., the central pixel 110 shown in Figure 9, or it may be derived from the values of one or more pixels sampled from a local neighborhood of pixels including the pixel of interest. In the preferred embodiment, the color gain determiner 135 finds the value of the boost $B_c(x,y)$ by applying a 2 dimensional LUT to the values $gm(x,y)$ and $ill(x,y)$. Typical values of $B_c(x,y)$ range from 0 to 5. The LUT should be a smoothly varying function of the 2 input variables. An example of such functions is described in European Patent Specification 0441558B1, granted January 7, 1998 and entitled "Improvements Relating to Control Data Arrays", which is incorporated herein by reference. Note that the color gain determiner 135 may alternatively be a function of one or more variables related to color (rather than 2 as described in the preferred embodiment.) Typically, the color gain determiner 135 will be a function of 1, 2 or 3 variables, and may be applied as a multi-dimensional LUT. As an alternative embodiment, the color signals input to the color gain determiner 135 need not be equivalent to the chrominance channels output from the luminance/chrominance converter 10. For instance, the red, green, and blue color channels could be input to the color gain determiner 135 while the green-magenta and illuminant channels are output from the luminance/chrominance converter 10. Those skilled in the art will recognize that similar results may be achieved by applying a color transformation to the original color digital image preceding input of color channels into the color gain determiner 135, and such a transformation does not significantly deviate from the description of the preferred embodiment.

[0068] In the preferred embodiment, the color gain determiner 135 applies a two dimensional LUT to the chrominance channels $gm(x,y)$ and $ill(x,y)$. The multi dimensional LUT may be formed using multivariate Gaussian functions. Each multivariate Gaussian function will be referred to as a color region pole, and may be described by a mean and a covariance about the mean. Each color region pole represents a region in color space for which the texture sharpening characteristics are preferentially treated. A color region pole is represented by a mean vector m and a covariance matrix k , as are well known in the art of statistics. In the preferred embodiment, for color region pole 0, the mean vector m_0 is represented as

$$m_0 = \begin{bmatrix} gm_0 \\ ill_0 \end{bmatrix}.$$

The covariance matrix of a multidimensional function is well known in the art and for color region pole 0,

$$k_0 = \begin{bmatrix} \sigma_{g^{m_0}}^2 & \sigma_{g^{m_0}} \sigma_{ill_0} \\ \sigma_{g^{m_0}} \sigma_{ill_0} & \sigma_{ill_0}^2 \end{bmatrix}$$

5
10 In the preferred embodiment, the mean vector m and the covariance matrix may be determined by a training procedure wherein an operator selects pixels of interest, having common material properties (for instance sky, skin, or grass) for one or more images. The mean and covariance of these selected pixels determine the mean m and covariance k of the color region pole. The actual numerical values of the mean vector m and the covariance matrix k are dependent upon the color and/or logarithmic transformations of the entire imaging chain. In the preferred embodiment, three color region poles are defined. Color region pole 0 represents skin. In this case,

$$m_0 = \begin{bmatrix} -26 \\ -109 \end{bmatrix}$$

20 and

$$k_0 = \begin{bmatrix} 2000 & 400 \\ 400 & 550 \end{bmatrix}$$

25 Color region pole 1 represents sky. In this case,

$$m_1 = \begin{bmatrix} -50 \\ 200 \end{bmatrix}$$

35 and

$$k_1 = \begin{bmatrix} 400 & 20 \\ 20 & 1000 \end{bmatrix}$$

40 Color region pole 2 represents grass. In this case,

$$m_2 = \begin{bmatrix} 153 \\ -69 \end{bmatrix}$$

50 and

$$k_2 = \begin{bmatrix} 800 & 0 \\ 0 & 600 \end{bmatrix}$$

55 [0069] Thus the color gain determiner 135 determines a value $Bc(x,y)$ by evaluating the following expression:

$$B_c(x, y) = \mu + \sum_{i=0}^P \mu_i \exp \left[-\frac{1}{2} (w - m_i)^T k_i (w - m_i) \right]$$

where P is the number of color region poles (in the preferred embodiment, $P = 3$), w is a vector representing the color of the local region of the image. In the preferred embodiment,

$$w = \begin{bmatrix} gm(x, y) \\ ill(x, y) \end{bmatrix}.$$

Note that the operation $(w - m_i)$ represents a color distance. Note that the value of w may be computed using only the selected pixel undergoing processing. Alternatively, the value of w may be determined by using the selected pixel plus some number of surrounding pixels, for example by using an averaging procedure, as is well known in the art of image processing. Thus, the value of w may be determined either as a point operation or as a local neighborhood operation. Also, μ is a base value of the color gain. In the preferred embodiment, $\mu = 1$. Finally, μ_i is the preferred ratio of the sharpness of color region pole i . Color region poles (such as skin and sky) for which less sharpening is desired have values of μ_i less than 0. Color region poles (such as grass) for which more sharpening is desired have values of μ_i more than 0. In the preferred embodiment, the values of μ_0 , μ_1 , and μ_2 are -0.5, -0.7 and 0.65 respectively. Thus, the value of B_c , the color weighting factor or color gain factor, in the preferred embodiment ranges from approximately 0.3 to 1.65.

[0070] Typically, the color gain determiner 135 is implemented as a LUT. Fig 15B shows an example of such a LUT. Note that the output value of this LUT can assume more than 2 values, and the LUT is said to be a near continuous function. Since the LUT is a near continuous function, the value of B_c is a non-binary value. Applying a multidimensional LUT is well known in the art of image processing, and is described for example in US Patent No. 4,500,919, which issued February 19, 1985 and is entitled "Color Reproduction System", and which is incorporated herein by reference. In this display, white represents a value of approximately 1.65 and black represents a value of approximately 0.3.

[0071] Note that although the color gain determiner 135 has herein been described as operating on chrominance channels corresponding to the original color digital image, it is possible that the chrominance channels may be modified with spatial filters or color transformations preceding the input of the chrominance channels to the color gain determiner 135. Moreover, although the color gain determiner 135 has herein been described as calculating an individual level of boost $B_c(x, y)$ to apply at each location (x, y) of the texture signal, dependent upon the local color as described by the chrominance channels $gm(x, y)$ and $ill(x, y)$, the individual level of boost $B_c(x, y)$ may instead, or in addition, be applied at each location (x, y) of the pedestal signal in order to further allow for independent control of enhancement applied to edges based on local color.

[0072] The SNR gain factor $B(x, y)$ output by the SNR gain determiner 134, the color gain factor $B_c(x, y)$ output from the color gain determiner 135, and the texture signal $n_{txt}(x, y)$ are then passed to a texture booster 136. The purpose of the texture booster 136 is to multiply each value of the texture signal $n_{txt}(x, y)$ by its associated gain. The output of the texture booster 136 is the enhanced texture signal $n'_{txt}(x, y)$, given by the equation:

$$n'_{txt}(x, y) = n_{txt}(x, y) * [(B(x, y) - B_m) * B_c(x, y) + B_m]$$

where B_m is a coefficient representing a minimum boost factor, which is 1 in the preferred embodiment. The effect, consequently, is that the color weighting or gain factor B_c causes an amplification (or a reduction) of high spatial frequency information for pixels that are closer in color to certain one(s) of the identified color regions, and a reduction of high spatial frequency information for pixels that are closer in color to certain other one(s) of the identified color regions. In other words, the color weighting factor $B_c(x, y)$ causes more (or less) noise reduction for pixels of interest that are closer to certain one(s) of the identified color regions.

[0073] Note that those skilled in the art will recognize that there are many equations by which a color gain factor $B_c(x, y)$ may be used with a SNR gain factor $B(x, y)$ to enhance a texture signal (or signal comprising mostly of detail.) This enhanced texture signal $n'_{txt}(x, y)$ output from the texture modifier 70 is then passed to the adder 80 in Fig. 3 (or to the avoidance adder 81 in the case of the alternative embodiment shown in Fig. 4.) Note when $(B(x, y) - B_m) * B_c(x, y) + B_m < 1$, the magnitude of the enhanced signal is less than the magnitude of the original texture signal $|n_{txt}(x, y)| > |n'_{txt}(x, y)|$. Thus, if the present invention is tuned such that in noisy regions, $(B(x, y) - B_m) * B_c(x, y) + B_m < 1.0$, the processed

digital image channel will appear to have less noise than the digital image channel input to the preference sharpener 2.

[0074] A detailed block diagram of the local SNR estimator 160 may be seen in Fig 12. The digital image channel $n(x,y)$ is first input to a local variance computer 170. The purpose of the local variance computer 170 is to attain a local estimate of activity. In the preferred embodiment, the local variance computer 170 calculates the local variance $\sigma_R^2(x,y)$ over a window of size 7 by 7 pixels, although it has been determined that alternately sized windows perform the task with nearly equivalent results.

[0075] The computation of the local variance of a set of values is well known in the art, and may be expressed as:

$$\sigma_R^2(x,y) = \Sigma p(x+i,y+j)^2/49 - (\Sigma p(x+i,y+j)/49)^2$$

Where i and j both range from -3 to +3, inclusive, for a total of 49 pixels included in the computation of the variance $\sigma_R^2(x,y)$. Calculation of the variance at edges involves mirroring, a technique well known in the art of image processing. The output of the local variance computer 170 is a local variance for each pixel in the digital image channel.

[0076] A subtractor 172 inputs both the digital image channel $n(x,y)$ and the texture signal $n_{txt}(x,y)$. The output of the subtractor 172 is the pedestal signal $n_{ped}(x,y)$. The subtractor 172 determines the pedestal signal by subtracting the texture signal $n_{txt}(x,y)$ from the digital image channel $n(x,y)$, according to the formula:

$$n_{ped}(x,y) = n(x,y) - n_{txt}(x,y)$$

Note that this pedestal signal $n_{ped}(x,y)$ is identical to the pedestal signal determined by the pedestal splitter 50. In a practical implementation of the present invention, the pedestal signal need only be computed once.

[0077] The noise table 6 and the pedestal signal $n_{ped}(x,y)$ are input an expected noise determiner 190. The noise table contains the relationship between expected noise $\sigma_N(i)$ and intensity i of the digital image channel $n(x,y)$. In the preferred embodiment, the noise table is a two column table, a sample of which is shown hereinbelow in a two column list. The first column represents intensity level and the second column represents the expected standard deviation of noise for that intensity level.

I	$\sigma_N(i)$
0	11.5
100	11.5
450	11.9
550	12.0
650	12.1
750	12.2
850	12.3
950	12.4
1050	12.3
1150	12.3
1250	12.3
1350	12.4
1450	12.4
1550	12.2
1650	11.9
1750	11.5
1850	11.4
1950	11.6
2050	12.1

(continued)

i	$\sigma_N(i)$
2150	12.8
2250	13.8
2350	15.0
2450	16.4
2550	17.9
2650	19.9

[0078] The noise table may also be specified in terms of a four column table, where the first column is the intensity i , and the second to fourth columns correspond to the standard deviations of noise expected in the red, green, and blue channels, respectively, of the digital image channel input to the preference sharpener 2. In this case, the expected standard deviation of noise for intensity level i of the luminance channel $n(x,y)$ is dependent upon the coefficients of the color rotation matrix used to generate the luminance channel, as given by:

$$\sigma_n(i) = \sqrt{(\sigma_r(i) \cdot rcoef)^2 + (\sigma_g(i) \cdot gcoef)^2 + (\sigma_b(i) \cdot bcoef)^2}$$

where $\sigma_n(i)$ is the expected standard deviation of noise at intensity i for the luminance digital image channel $n(x,y)$. Likewise, $\sigma_r(i)$, $\sigma_g(i)$, $\sigma_b(i)$ are the expected standard deviations of noise of the respective red, green, and blue digital image channels at intensity i , and $rcoef$, $gcoef$, and $bcoef$ are the coefficients of the color rotation matrix used to generate the luminance digital image channel $n(x,y)$. As previously mentioned, in the preferred embodiment all three of these coefficients equal 1/3. Consequently, the luminance channel given in the preferred embodiment is given by:

$$\sigma_n(i) = 1/3 \cdot \sqrt{\sigma_r(i)^2 + \sigma_g(i)^2 + \sigma_b(i)^2}$$

Thus a noise table corresponding to the luminance digital image channel $n(x,y)$ may be either supplied or calculated by the expected noise determiner 190.

[0079] The expected noise determiner 190 outputs an intensity dependent estimate of the standard deviation of noise $\sigma_N(x,y)$ for the intensity level of the pedestal signal $n_{ped}(x,y)$. The expected noise determiner 190 either performs a table look-up, which relates the intensity of the pedestal signal to the expected standard deviation of noise $\sigma_N(x,y)$, or calculates the expected noise, as described above. The output $\sigma_N(x,y)$ of the expected noise determiner 190 may be represented as:

$$\sigma_N(x,y) = \sigma_N(n_{ped}(x,y)).$$

If the value of $n_{ped}(x,y)$ is not coincident with a value of i for which $\sigma_n(i)$ exists, then linear interpolation is used to determine a value for $\sigma_n(n_{ped}(x,y))$. Linear interpolation is well known in the art and will not be further discussed. A plot of a noise table can be seen in Fig. 14.

[0080] The procedure for generating a noise table such as the one illustrated in Fig. 14 can be found in U.S. Patent No. 5,641,596 issued June 24, 1997 to Gray et al. or in U.S. Patent No. 5,923,775 issued July 13, 1999 to Snyder et al., both of which are incorporated herein by reference.

[0081] An SNR calculator 180 inputs both the expected standard deviation of noise $\sigma_N(x,y)$ output from the expected noise determiner 190, and the local variance $\sigma_R^2(x,y)$ output from the local variance computer 170. The SNR calculator 180 estimates a local signal to noise ratio $SNR(x,y)$ by computing the ratio of the signal variance (assuming that the local variance $\sigma_R^2(x,y)$ is the sum {in an RMS sense} of the variance due to signal and the variance due to noise.) Thus, the output of the SNR calculator 180 is an estimate of the local SNR at each pixel location, according to the following formula:

$$SNR(x,y) = \sqrt{(\sigma_R^2(x,y) - \sigma_N^2(x,y)) / \sigma_N^2(x,y)}$$

Note that before attempting the square root operation, the value of $\sigma_R^2(x,y) - \sigma_N^2(x,y)$ is checked to ensure that it is positive. If it is not, the value of SNR (x,y) is set to equal 0.0.

[0082] Notice that many other estimates of the signal to noise ratio may be formulated. For example, an alternative embodiment of the local SNR estimator 160 is shown in Fig. 13. The purpose of this embodiment is to provide an estimate of the local SNR by a means that is computationally less expensive than that shown in Fig. 12. In this alternative embodiment, the noise and signal deviation are both assumed to be Gaussian distributed, which allows for the use of mean absolute deviation (MAD) rather than standard deviation. The calculation of mean absolute deviation does not involve the computationally expensive square root and squaring operations. In this embodiment, the texture signal $n_{\text{txt}}(x,y)$ and the digital image channel $n(x,y)$ are again input to the subtractor 172 for the purpose of producing the pedestal signal $n_{\text{ped}}(x,y)$, as previously described. The pedestal signal $n_{\text{ped}}(x,y)$ will again be used as an input to the expected noise determiner 190 to determine the value of $\sigma_N(x,y)$.

[0083] The texture signal $n_{\text{txt}}(x,y)$ is input to the local MAD calculator 200 in order to determine the local mean absolute deviation in the texture signal. By assuming that the mean value of the texture signal $n_{\text{txt}}(x,y)$ is 0, the formula for computing the local MAD used by the local MAD calculator 200 is given as:

$$\text{MAD}_R(x,y) = \sum |n_{\text{txt}}(x+i,y+j)| / 49$$

where i and j preferably both range from -3 to +3, for a total of 49 pixels included in the calculation of the mean absolute deviation. The local MAD calculator 200 outputs the value of the mean absolute deviation $\text{MAD}_R(x,y)$, which will later be used in calculations necessary for determining the estimate of the local SNR.

[0084] The purpose of a noise table adjuster 210 is to compensate for the differences between calculating $\text{MAD}_R(x,y)$ of the texture signal, and calculating the standard deviation $\sigma_R(x,y)$ of the digital image channel. The noise table adjuster 210 modifies the luminance noise table for the purpose of using mean absolute deviation, rather than standard deviation. Assuming, as previously mentioned, that the luminance noise table contains the relationship between expected noise $\sigma_n(i)$ and intensity i , it is known in the art that the MAD of a Gaussian distribution is related to the standard deviation σ of that distribution by the following relationship:

$$\text{MAD} = \text{sqrt}(2/\pi) \sigma$$

Thus, in order to translate the luminance noise table from the metric of standard deviation to the metric of MAD, each value of $\sigma_n(i)$ must be scaled by about 0.8.

[0085] In addition, there is the difference between this alternative embodiment and the preferred embodiment concerning which signal is used for computation of the local activity measure. In the preferred embodiment, the variance of the digital image channel is computed. In this alternative embodiment, the texture signal is used to calculate the value of $\text{MAD}_R(x,y)$, providing a computational benefit when the signal mean is assumed to be zero. Thus the relationship of the texture signal $n_{\text{txt}}(x,y)$ to the digital image channel $n(x,y)$ must be understood in order to adjust the luminance noise table to account for this. This relationship can be easily expressed only if the digital filtering process used to generate the texture signal is linear. For the sake of this step, the effects of the artifact avoidance signal are ignored, and the texture signal is assumed to be equivalent to the highpass signal. For a general $(n \times m)$ filter with coefficients $h(i,j)$, wherein $i = -(n-1)/2, -(n-3)/2, \dots, n-1/2$ and $j = -(m-1)/2, -(m-3)/2, \dots, m-1/2$, the variance of the filtered signal is related to the variance of the original signal by the following relationship:

$$\sigma_{\text{fs}}^2 = \sigma_{\text{os}}^2 \sum \sum h^2(i,j)$$

where σ_{fs}^2 represents the variance of the filtered signal, σ_{os}^2 represents the variance of the original signal, and the summation occurs over i and j .

[0086] Thus, in the case of the present invention, the filter $h(i,j)$ is assumed to be expressed as:

$$h(i,j) = \delta(i,j) - g(i,j)$$

where $\delta(i,j)$ is the delta function which has a value of one at $i = 0$ and $j = 0$ and a value of zero everywhere else. This expression accounts for the previously described process of generating the texture signal (ignoring the artifact avoidance signal) by blurring the digital image channel with a Gaussian filter $g(i,j)$, then subtracting the result from the digital image channel. Thus, considering this description, the noise table adjuster 210 must scale each value of the luminance

noise table by a factor to account for the estimation of local activity on the texture signal rather than the digital image channel. This factor f can be expressed as:

$$f = \text{sqrt} (\sum \sum h(i,j)^2)$$

where the summation occurs over i and j , and $h(i,j)$ is as previously described. The square root of the result is taken because the luminance noise table is in the units of standard deviation, rather than variance.

[0087] Thus, the operation of the noise table adjuster 210 is to scale each entry $\sigma_n(i)$ of the luminance noise table by a factor m equal to

$$m = f \cdot \text{sqrt} (2/\pi)$$

in order to compensate for the calculation of MAD rather than standard deviation, and the calculation with the texture signal rather than the digital image channel.

[0088] The modified noise table output from the noise table adjuster 210 is passed as an input to the expected noise determiner 190 for the purpose of attaining an estimate of the expected MAD from noise, $MAD_N(x,y)$. The expected noise determiner 190 uses the modified noise table to determine the value of $MAD_N(x,y)$ for each value of $n_{ped}(x,y)$. The output of the expected noise determiner 190 is the value of $MAD_N(x,y)$ at each location of the digital image channel.

[0089] The SNR calculator 180 of the alternative embodiment shown in Fig. 13 acts essentially the same as the SNR calculator 180 of the preferred embodiment. The purpose of the SNR calculator 180 is to output a value of estimated SNR at each location of the digital image channel. To this end, the SNR calculator 180 calculates the following formula:

$$\text{SNR}(x,y) = \text{sqrt}(MAD_R(x,y)^2 - MAD_N(x,y)^2) / MAD_N(x,y)$$

The estimate of the local SNR is output from the local SNR estimator 160 and input to the gain determiner 134, shown in Fig. 11 in order to allow for the determination of the appropriate boost factor for the texture signal at each location.

[0090] In summary, the use of the pedestal splitter 50, decomposing the image into a signal containing the image detail and another signal primarily containing the large image edges, allows for independent control of the enhancement applied to edges and texture. Edges are reshaped by a morphological operation in order to increase edge contrast and avoid boundary artifacts. Texture is enhanced with regard to an estimate of the local signal to noise ratio and, in accordance with the invention, with regard to an estimate of local color.

[0091] The present invention has been described with reference to a preferred embodiment. Changes may be made to the preferred embodiment without deviating from the scope of the present invention. For instance, in the preferred embodiment, a calculation of the local SNR is made for each pixel of the digital image channel. However, it is conceivable that the calculation of the SNR could be made for only every N^{th} pixel in order to reduce computational cost, and replicated or interpolated. Calculation of the local SNR was described with reference to a local variance calculation on the digital image channel $n(x,y)$. Additionally, an alternative embodiment described a method by which local SNR may be estimated based upon using a MAD calculation upon the texture signal $n_{tx}(x,y)$. Those skilled in the art will recognize that the local SNR may also be estimated by performing a variance calculation on the texture signal, or by performing a MAD calculation upon the digital image channel $n(x,y)$. Such modifications to the preferred embodiment do not significantly deviate from the scope of the present invention.

Claims

1. A digital image processing method for enhancing a color digital image composed of a plurality of pixels having color values representative of the image, said method comprising the steps of:

- a) identifying a local neighborhood of pixels including a pixel of interest;
- b) using the color values of one or more of the pixels from the local neighborhood to calculate a color weighting factor for the pixel of interest, the color weighting factor being a near continuous function of the location of the one or more pixels from the local neighborhood in a color space;
- c) applying a spatial filter to the values of pixels from the local neighborhood to calculate a first signal value and a second signal value for the pixel of interest;
- d) using the color weighting factor to modify either the first value or the second value;

e) generating an enhanced pixel value for the pixel of interest from the first value and second values; and
f) repeating steps (a) through (e) for other pixels in the color digital image.

2. The method claimed in claim 1, wherein the spatial filter is a function of the pixel values of the local neighborhood.

3. The method claimed in claim 1, wherein the spatial filter is a non-linear function of the pixel values of the local neighborhood.

4. The method claimed in claim 1, wherein the first signal value predominantly relates to the low spatial frequency information and the second signal value predominantly relates to the high spatial frequency information.

5. The method claimed in claim 4, wherein step c) further comprises the steps of applying a spatial filter to the first signal values for pixels in the color digital image to calculate a gradient signal value from the first signal values, and using the gradient signal value to modify either the first signal value or second signal value.

6. The method as claimed in claim 5 wherein the gradient signal values are modified with a single valued function.

7. The method claimed in claim 1, wherein step b) further comprises the steps of identifying one or more color regions in the color space, calculating a pixel color coordinate in the color space for the pixel of interest, and calculating the color weighting factor as a near continuous function of the pixel color coordinate and the one or more color regions.

8. The method claimed in claim 1, further comprising the steps of using the values of pixels sampled from the local neighborhood to calculate a statistical weighting factor for the pixel of interest and using the statistical weighting factor to modify either the first signal value or second signal value.

9. The method claimed in claim 8, further comprising the steps of identifying a noise parameter relating to the expected noise level for the pixel of interest and using the statistical weighting factor and the noise parameter to modify either the first signal value or second signal value.

10. The method claimed in claim 9, wherein the step of identifying a noise parameter comprises the steps of calculating an intensity signal value for the pixel of interest and identifying a noise parameter relating to the expected noise level for the pixel of interest, the noise parameter being a function of the intensity signal value.

11. The method claimed in claim 7, wherein step b) further comprises the steps of calculating a neighborhood color value derived from the values of one or more pixels sampled from a local neighborhood of pixels including the pixel of interest, calculating color distance(s) from the neighborhood color value to the identified color region(s) and using the color distance values to calculate the color weighting factor.

12. The method claimed in claim 11, wherein the neighborhood color value is derived only from the pixel of interest.

13. The method claimed in claim 11, wherein the color weighting factor is calculated as a Gaussian function of the color distance(s).

14. The method claimed in claim 13, wherein the color weighting factor is calculated using a Gaussian function of the color distance values for each identified color region.

15. The method claimed in claim 14, wherein the Gaussian functions are combined in an additive manner to calculate the color weighting factor.

16. The method claimed in claim 14, wherein the Gaussian functions are combined in a multiplicative manner to calculate the color weighting factor.

17. The method claimed in claim 7, wherein the color weighting factor causes an amplification of high spatial frequency information for pixels of interest that are closer in color to one of the identified color regions.

18. The method claimed in claim 7, wherein the color weighting factor causes a reduction of high spatial frequency information for pixels of interest that are closer in color to one of the identified color regions.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5

FIG. 5A
(ALTERNATIVE EMBODIMENT OF
THE PEDESTAL SPLITTER)





FIG. 5B
(ALTERNATIVE EMBODIMENT OF
THE PEDESTAL SPLITTER)






FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10




FIG. 11




FIG. 12




FIG. 13




FIG. 14




FIG. 15A





Fig. 15B

Best Available Copy

THIS PAGE BLANK (USPTO,



Europäisches Patentamt

(19) European Patent Office

Office européen des brevets



(11)

EP 1 209 621 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
10.09.2003 Bulletin 2003/37

(51) Int Cl.7: G06T 5/20, H04N 1/409

(43) Date of publication A2:
29.05.2002 Bulletin 2002/22

(21) Application number: 01203319.7

(22) Date of filing: 03.09.2001

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• Gallagher, Andrew C.,
c/o Eastman Kodak Company
Rochester, New York 14650-2201 (US)
• Gindele, Edward B.,
c/o Eastman Kodak Company
Rochester, New York 14650-2201 (US)

(30) Priority: 13.09.2000 US 661268

(71) Applicant: EASTMAN KODAK COMPANY
Rochester, New York 14650 (US)

(74) Representative: Haile, Helen Cynthia et al
Kodak Limited
Patent, W92-3A,
Headstone Drive
Harrow, Middlesex HA1 4TY (GB)

(54) A method for enhancing a digital image based upon pixel color

(57) A method is described for enhancing a color digital image by first identifying a local neighborhood of pixels including a pixel of interest and then using the color values of one or more of the pixels from the local neighborhood to calculate a color weighting factor for the pixel of interest, where the color weighting factor is a near continuous function of the location of the pixel of interest in a color space. A spatial filter is applied to the

values of pixels from the local neighborhood to calculate a first signal value and a second signal value for the pixel of interest. The color weighting factor is then used to modify either the first value or the second value, and an enhanced pixel value is then generated for the pixel of interest from the first value and second values. The process is then repeated for other pixels in the color digital image.




FIG. II



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 20 3319

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
A	EP 0 971 314 A (EASTMAN KODAK CO) 12 January 2000 (2000-01-12) * paragraph [0013] - paragraph [0014] * * paragraph [0022] - paragraph [0023] * ---	1-18	G06T5/20 H04N1/409
D,A	EP 0 441 558 A (CROSFIELD ELECTRONICS LTD) 14 August 1991 (1991-08-14) * page 3, line 4 - line 34 * * page 5, line 39 - line 43 * -----	1-18	
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
			H04N G06T
The present search report has been drawn up for all claims			
Place of search MUNICH		Date of completion of the search 11 July 2003	Examiner Rockinger, O
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons</p> <p>& : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03.02 (P04C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 01 20 3319

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

11-07-2003

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
---EP 0971314--- A	12-01-2000	US 6317521 B1	13-11-2001
		EP 0971314 A2	12-01-2000
		JP 2000075852 A	14-03-2000

EP 0441558 A	14-08-1991	DE 69128548 D1	12-02-1998
		DE 69128548 T2	16-04-1998
		EP 0441558 A1	14-08-1991
		JP 2969010 B2	02-11-1999
		JP 4227585 A	17-08-1992
		US 5105469 A	14-04-1992

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

THIS PAGE BLANK (USPTO)